HOUSTONKEMP Economists

Memo

To	Greg Houston
From	Simon Wheatley
Copy	Brendan Quach, Nick Twort
Subject	RBA discussion paper 2019-04
Date	10 March 2022

You will recall that in August 2021, I wrote a report, *A new MRP estimate? A review of RBA discussion paper 2019-04*, for Energy Networks Australia (ENA) that was subsequently submitted to the Australian Energy Regulator (AER). The report examines the way in which a recent Reserve Bank of Australia (RBA) paper,

> Mathews, Thomas, *A history of Australian equities*, RBA discussion paper RDP 2019-04,

produces estimates of the long-run market risk premium (MRP) for Australia and concludes that the estimates will be downwardly biased. In fact, the report shows that not only should Mathews' estimates of the MRP be downwardly biased but that one should expect the market return that he produces to be downwardly biased in every year.

A puzzling feature of the market returns that Mathews produces, however, is that while the returns are on average lower than their counterparts drawn from the series that Brailsford, Handley and Maheswaran (2008, 2012) and NERA (2013, 2015) construct and we extend, the returns in some years are higher. [1] On further investigation I have discovered that this puzzling feature is almost surely caused by an error in the program that Mathews uses to generate returns.

The error is in an ifelse statement and the impact of the error is that Mathews will have gotten the capital change classification scheme that he uses back to front. So, for example, Mathews will have ended up classifying most share issuances as stock splits, almost all stock splits as share issuances and buybacks as reverse splits. It is not possible to say whether the error will lead the returns that Mathews produces to be higher or lower but the existence of the error must cast doubt on the reliability of his work.

Capital changes

Mathews computes a series of with-dividend returns from 1917 to 1979 to the largest 100 Australian stocks by market capitalisation, using dividends, prices and shares outstanding for each stock and for each quarter taken from various issues of the Sydney Stock Exchange *Official Gazette*, and then links these returns to series provided by Datastream from 1980 to 2019.

The problem that Mathews faces – and acknowledges that he faces – in using dividends, prices and shares outstanding alone to compute a series of returns is how to correctly handle capital changes. Capital changes include bonus issues, stock splits, reverse splits, the issuance of new shares and share buybacks.

[1] Brailsford, T., J. Handley and K. Maheswaran, *Re-examination of the historical equity risk premium in Australia*, Accounting and Finance 48, 2008, pp. 73-97.

Brailsford, T., J. Handley and K. Maheswaran, *The historical equity risk premium in Australia: Post-GFC and 128 years of data*, Accounting and Finance, 2012, pp. 237-247.

NERA, *Market, size and value premiums: A report for the ENA*, June 2013.

NERA, *The market risk premium: Analysis in response to the AER's Draft Rate of Return Guidelines: A report for the Energy Networks Association*, October 2013.

NERA, *Further assessment of the historical MRP: Response to the AER's final decisions for the NSW and ACT electricity distributors: A report for ActewAGL Distribution, AGN, APA, AusNet Services, CitiPower, Energex, Ergon Energy, Jemena Electricity Networks, Powercor, SA Power Networks and United Energy*, June 2015.

The way in which Mathews says that he addresses the problem of handling capital changes is to use changes in the number of shares outstanding as a signal that capital changes have taken place and a crude classification scheme to adjust for the changes. Mathews says in RDP 2019-04 that he assumes that: [2]

- if the number of shares that a firm has outstanding more than doubles over a quarter, the firm has carried out a stock split; and
- if the number of shares that a firm has outstanding does not more than double over a quarter, changes in the number of shares outstanding are the result of the firm buying back or issuing shares at a fair price.

This classification scheme, as my August 2021 report makes clear, will misclassify a number of capital changes. The scheme will misclassify:

- bonus issues and stock splits that a firm carries out that do not lead to a more than doubling of the number of shares of stock that the firm has outstanding;
- reverse splits; and
- issues of new shares at a fair price that involve a more than doubling of the number of shares that a firm has outstanding.

It is likely that most stock splits that a firm executes involve a more than doubling of the number of shares the firm has outstanding and so the risk of the scheme misclassifying stock splits as share issuances is small. The classification scheme will misclassify reverse splits but reverse splits are rare and tend to involve distressed and so low-market-capitalisation firms. Consequently, the impact of the scheme misclassifying reverse splits will be small. It is also rare that firms issue new shares at a fair price and more than double the number of shares that they have outstanding. Moreover, any firms doing so are likely to be low-market-capitalisation firms and so the impact of misclassifying issuances of new shares as stock splits will also be small.

Bonus issues, though, are and were common and the classification scheme that Mathews says that he uses will misclassify virtually all of them as issuances of new shares. Bonus issues are accompanied by stock price declines that are a purely mechanical consequence of the issues and in no way harm investors. The classification scheme, that Mathews says that he uses, will feed what appear to be the large negative returns associated with the stock price declines into any series of returns constructed using the scheme. As a result, one would expect not only Mathews' estimates of the MRP to be downwardly biased but the market return that he produces to be downwardly biased in every year. I find, however, that while the market returns that Mathews produces are on average lower than their Brailsford, Handley and Maheswaran (2008, 2012) and NERA (2013, 2015) counterparts, the returns in some years are higher. [3] This puzzling behaviour prompted me to take a closer look at Mathews' code and this closer look in turn led me to uncover an error in the function that Mathews uses to generate returns. The coding error in this function almost surely explains why

[2] Mathews, T., *A history of Australian equities*, RBA research discussion paper RDP 2019-04, p 32.

[3] Brailsford, T., J. Handley and K. Maheswaran, *Re-examination of the historical equity risk premium in Australia*, Accounting and Finance 48, 2008, pp. 73-97.

Brailsford, T., J. Handley and K. Maheswaran, *The historical equity risk premium in Australia: Post-GFC and 128 years of data*, Accounting and Finance, 2012, pp. 237-247.

NERA, *Market, size and value premiums: A report for the ENA*, June 2013.

NERA, *The market risk premium: Analysis in response to the AER's Draft Rate of Return Guidelines: A report for the Energy Networks Association*, October 2013.

NERA, *Further assessment of the historical MRP: Response to the AER's final decisions for the NSW and ACT electricity distributors: A report for ActewAGL Distribution, AGN, APA, AusNet Services, CitiPower, Energex, Ergon Energy, Jemena Electricity Networks, Powercor, SA Power Networks and United Energy*, June 2015.

the returns that Mathews produces are in some years higher than their Brailsford, Handley and Maheswaran and NERA counterparts.

Coding error

Mathews uses the language R to write the programs that he employs and the ifelse statement to operationalise the classification scheme that is supposed to handle capital changes. RDocumentation states that:[4]

> Description
>
> 'ifelse' returns a value with the same shape as 'test' which is filled with elements selected from either 'yes' or 'no' depending on whether the element of 'test' is 'TRUE' or FALSE'.
>
> Usage
>
> ifelse(test, yes, no)

In the ifelse statement that Mathews uses to operationalise the classification scheme, he gets the 'yes' and 'no' elements around the wrong way. In other words, Mathews writes the ifelse statement as ifelse(test, no, yes) rather than as ifelse(test, yes, no). The appendix to this memo contains both the function that Mathews uses to compute returns, taken from the RBA's web site, and a corrected version of the function.

Mathews' code also uses a different cut-off to determine whether a share issuance or a stock split has taken place. Mathews states in RDP 2019-04 that he presumes that when the number of shares that a firm has outstanding more than doubles over a quarter, the firm has carried out a stock split. So in RDP 2019-04, Mathews says that he uses a cut-off of a 100 per cent increase in the number of shares outstanding. His program, however, uses instead a cut-off of a 90 per cent increase in the number of shares outstanding.

Thus while Mathews states on page 32 of RBA 2019-04 that he computes the without-dividend return to a portfolio of the largest 100 stocks as: [5]

$$\frac{\sum_{\text{all } j} p_{jt} q_{jt}}{\sum_{\text{all } j \text{ s.t. } q_{jt} \le 2q_{jt-1}} p_{jt-1} q_{jt} + \sum_{\text{all } j \text{ s.t. } q_{jt} > 2q_{jt-1}} p_{jt-1} q_{jt-1}} - 1$$

where:

- p_{jt} is the price of stock *j* at time *t*; and
- q_{jt} is the number of shares outstanding of stock *j* at time *t*,

in practice, Mathews computes the without-dividend return as:

[4] rdocumentation.org/packages/base/versions/3.6.2/topics/ifelse

[5] I presume that the statement at the bottom of page 32 that a split is presumed to have occurred when $q_{jt} > 100 q_{jt-1}$ is a mistake.

$$\frac{\sum_{\text{all } j} p_{jt} q_{jt}}{\sum_{\text{all } j \text{ s.t. } q_{jt} > 1.9 q_{jt-1}} p_{jt-1} q_{jt} + \sum_{\text{all } j \text{ s.t. } q_{jt} \leq 1.9 q_{jt-1}} p_{jt-1} q_{jt-1}} - 1$$

In other words, Mathews assumes that:

- if the number of shares that a firm has outstanding rises by more than 90 per cent over a quarter, changes in the number of shares outstanding are the result of the firm issuing shares at a fair price; and
- if the number of shares that a firm has outstanding does not rise by more than 90 per cent over a quarter, changes in the number of shares outstanding are the result of a stock split or reverse split.

This alternative scheme will also misclassify a number of capital changes. The scheme will misclassify:

- bonus issues and stock splits that a firm carries out that lead to the number of shares of stock that the firm has outstanding rising by more than 90 per cent;
- share buybacks; and
- issues of new shares at a fair price that do not involve the number of shares of stock that the firm has outstanding rising by more than 90 per cent.

Few bonus issues involve the number of shares that a firm has outstanding rising by more than 90 per cent. So the alternative scheme will in general correctly classify bonus issues. On the other hand, almost all stock splits that a firm carries out involve the number of shares of stock that the firm has outstanding rising by more than 90 per cent and so the scheme will misclassify almost all stock splits as share issuances. Stock splits, like bonus issues, are accompanied by stock price declines that are a mechanical consequence of the splits and do not harm investors. The alternative scheme will feed what appear to be the large negative returns associated with the stock price declines into the series of returns that Mathews constructs.

The alternative scheme will also misclassify share buybacks as reverse splits. Share buybacks and reverse splits cause the number of shares of stock that a firm has outstanding to fall. Share buybacks at a fair price should have no impact, all else constant, on the prices of shares. Reverse splits, in contrast, are accompanied by stock price increases that are a mechanical consequence of the splits and do not benefit investors. The alternative scheme, because it will misclassify share buybacks as reverse splits, will adjust share prices around buybacks for the increases that one would expect to see were the buybacks to be reverse splits and, as a result, will construct returns around buybacks that are downwardly biased. The scheme will look around buybacks for the stock price increases that are a mechanical consequence of reverse splits, will not find them and will interpret the absence of the price increases that are a consequence of reverse splits as bad news.

Finally, the alternative scheme will misclassify issues of new shares at a fair price that do not involve the number of shares of stock that the firm has outstanding rising by more than 90 per cent as stock splits. Share issuances and stock splits cause the number of shares of stock that a firm has outstanding to rise. Share issues at a fair price should have no impact, all else constant, on the prices of shares. Stock splits, in contrast, are accompanied by stock price declines that are a mechanical consequence of the splits and do not harm investors. The alternative scheme, because it will misclassify issues of new shares at a fair price that do not involve the number of shares of stock that the firm has outstanding rising by more than 90 per cent as stock splits, will adjust share prices around the issues for the declines that one would expect to see were the issues to be stock splits and, as a result, will construct returns around the issuances that are upwardly biased. The scheme will look around the issuances for the stock price declines that are a mechanical consequence of stock splits, will not find them and will interpret the absence of the price declines that are a consequence of stock splits as good news.

This analysis indicates that it is not possible to say whether the coding error in the function that Mathews uses to construct returns will lead the returns that he produces to be higher or lower. The existence of the error, though, must cast doubt on the reliability of Mathews' work.

Illustration

It will be helpful to provide an illustration of how the alternative classification scheme that Mathews uses – inadvertently, one would presume – can generate errors in returns. The example that I employ uses an all-equity firm that has, to begin with, 100 shares outstanding. Each share trades initially at $1.00 and the value of an initial investment of $1.00 in one share does not change over the next 30 days. On day 10, however, the firm issues 25 new shares at the market price of $1.00 and then on day 20 the firm carries out a 4-for-1 stock split. The firm does not engage in any other capital changes or pay out any dividends over the 30 days.

Table 1 below provides:

- the number of the firm's shares outstanding;
- the firm's share price;
- the firm's value;
- the true value of an initial investment of $1.00 in one share; and
- the value of an initial investment of $1.00 in one share generated by the R function that Mathews uses .[6]

Figure 1 below plots the true value of an initial investment of $1.00 in one share and the value of an initial investment of $1.00 in one share generated by the R function that Mathews uses against time. Table 1 and Figure 1 show that the R function that Mathews uses:[7]

- overstates the return to a share in the firm around the share issuance; and
- understates the return to a share in the firm around the stock split.

In the example, the impact of the stock split on the return that one computes is greater than the impact of the share issuance on the return. This will in general be the case. Whether the error in Mathews' code leads the MRP to be over- or underestimated, however, will depend not only the relative impacts of the error on returns around buybacks, share issuances and stock splits, but also on the relative frequencies of these capital changes. While it is not possible to say whether the error in Mathews' code will lead the returns that he produces to be higher or lower, however, the existence of the error must cast doubt on the reliability of his work.

[6] Nick Twort, who as you know is an experienced R programmer, was kind enough to run Mathews' code and a corrected version of the code. Again, both versions of the code appear in the appendix to this memo.

[7] In contrast, largely because there are no bonus issues in the example, the classification scheme that Mathews says that he uses will produce the true value of an initial investment of $1.00 in one share at each date. In practice, bonus issues are common and so, as I make clear in my August 2021 report, the classification scheme that Mathews says that he uses will in general underestimate returns.

Table 1: Illustrative example

Date	Shares outstanding	Share price	Firm value	True value of initial investment of $1.00	Mathews value of initial investment of $1.00
1	100	1.00	100	1.00	1.00
2	100	1.00	100	1.00	1.00
3	100	1.00	100	1.00	1.00
4	100	1.00	100	1.00	1.00
5	100	1.00	100	1.00	1.00
6	100	1.00	100	1.00	1.00
7	100	1.00	100	1.00	1.00
8	100	1.00	100	1.00	1.00
9	100	1.00	100	1.00	1.00
10	125	1.00	125	1.00	1.25
11	125	1.00	125	1.00	1.25
12	125	1.00	125	1.00	1.25
13	125	1.00	125	1.00	1.25
14	125	1.00	125	1.00	1.25
15	125	1.00	125	1.00	1.25
16	125	1.00	125	1.00	1.25
17	125	1.00	125	1.00	1.25
18	125	1.00	125	1.00	1.25
19	125	1.00	125	1.00	1.25
20	500	0.25	125	1.00	0.31
21	500	0.25	125	1.00	0.31
22	500	0.25	125	1.00	0.31
23	500	0.25	125	1.00	0.31
24	500	0.25	125	1.00	0.31
25	500	0.25	125	1.00	0.31
26	500	0.25	125	1.00	0.31
27	500	0.25	125	1.00	0.31
28	500	0.25	125	1.00	0.31
29	500	0.25	125	1.00	0.31
30	500	0.25	125	1.00	0.31

Figure 1: Illustrative example

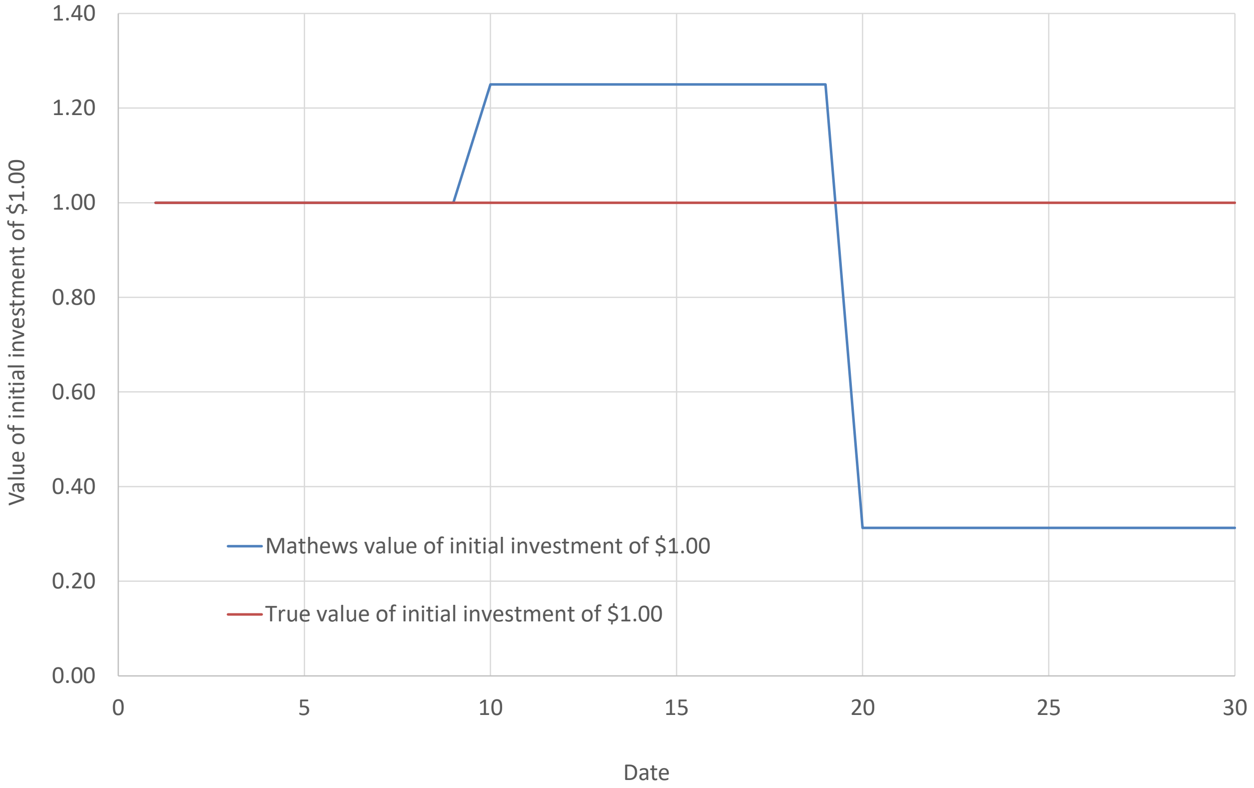

Appendix

The function that Mathews uses to compute returns appears below with the line of code that contains an error in red. The function was taken from the program ‘additional functions.R’ that can be found at rba.gov.au/publications/rdp/2019/2019-04/supplementary-information.html.

```
# calculate price indices from company-level data
udf_price_index <- function(data, sector) {
  # this function takes a given sectoral subset and calculates a price index
  # the methodology is explained in the paper, but the main trick is to make sure it is unaffected by
compositional change
  # that means include only companies that are in both periods

  if (sector == "Total") {
    sector <- c("Resources", "Financial", "Other")
  }

  data <- data %>%
    filter(Broad.sector %in% sector) %>%
    complete(Dates, nesting(Security, Broad.Company.Name)) %>% #make sure the lags are calculated
correctly when companies are moving in and out of the index
    group_by(Security) %>%
    arrange(Dates) %>%
    mutate(
      company_gr = Market.Capitalisation / lag(Market.Capitalisation, 1),
      share_gr = Shares / lag(Shares, 1),
      share_gr = ifelse(share_gr > 1.9, share_gr, 1),
      company_gr = company_gr / share_gr,
      weight = lag(Market.Capitalisation, 1),
      weight = ifelse(is.na(weight), 0, weight)
    ) %>%
    group_by(Dates) %>%
    summarise(
      index_gr = weighted.mean(company_gr, weight, na.rm = T),
      index_gr = ifelse(is.na(index_gr), 1, index_gr)
    ) %>%
    mutate(Price_index = cumprod(index_gr)) %>%
    select(-index_gr) %>%
    mutate(Broad.sector = ifelse(length(sector) == 3, "Total", sector))

  return(data)

}
```

The corrected version of the function appears below with the line of code that has been corrected in blue.

```
# calculate price indices from company-level data
udf_price_index_amended <- function(data, sector) {
  # this function takes a given sectoral subset and calculates a price index
  # the methodology is explained in the paper, but the main trick is to make sure it is unaffected by
compositional change
  # that means include only companies that are in both periods

  if (sector == "Total") {
    sector <- c("Resources", "Financial", "Other")
  }

  data <- data %>%
    filter(Broad.sector %in% sector) %>%
    complete(Dates, nesting(Security, Broad.Company.Name)) %>% #make sure the lags are calculated
correctly when companies are moving in and out of the index
    group_by(Security) %>%
    arrange(Dates) %>%
    mutate(
      company_gr = Market.Capitalisation / lag(Market.Capitalisation, 1),
      share_gr = Shares / lag(Shares, 1),
      share_gr = ifelse(share_gr > 2, 1, share_gr), # Adjustment on this line
      company_gr = company_gr / share_gr,
      weight = lag(Market.Capitalisation, 1),
      weight = ifelse(is.na(weight), 0, weight)
    ) %>%
    group_by(Dates) %>%
    summarise(
      index_gr = weighted.mean(company_gr, weight, na.rm = T),
      index_gr = ifelse(is.na(index_gr), 1, index_gr)
    ) %>%
    mutate(Price_index = cumprod(index_gr)) %>%
    select(-index_gr) %>%
    mutate(Broad.sector = ifelse(length(sector) == 3, "Total", sector))

  return(data)

}
```